



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

CUTTING PLANE GENERATION THROUGH SPARSE PRINCIPAL COMPONENT ANALYSIS

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Dey, S.S., Kazachkov, A., Lodi, A., Munoz, G. (2022). CUTTING PLANE GENERATION THROUGH SPARSE PRINCIPAL COMPONENT ANALYSIS. SIAM JOURNAL ON OPTIMIZATION, 32(2), 1319-1343 [10.1137/21M1399956].

Availability:

This version is available at: <https://hdl.handle.net/11585/905200> since: 2024-03-01

Published:

DOI: <http://doi.org/10.1137/21M1399956>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

CUTTING PLANE GENERATION THROUGH SPARSE PRINCIPAL COMPONENT ANALYSIS

SANTANU S. DEY*, ALEKSANDR M. KAZACHKOV†, ANDREA LODI‡, AND GONZALO MUÑOZ§

Abstract. Quadratically-constrained quadratic programs (QCQPs) are optimization models whose remarkable expressiveness have made them a cornerstone of methodological research for non-convex optimization problems. However, modern methods to solve a general QCQP fail to scale, encountering computational challenges even with just a few hundred variables. Specifically, a semi-definite programming (SDP) relaxation is typically employed, which provides strong dual bounds for QCQPs, but relies on memory-intensive algorithms. An appealing alternative is to replace the SDP with an easier-to-solve linear programming relaxation, while still achieving strong bounds. In this work, we make advances towards achieving this goal by developing a computationally-efficient linear cutting plane algorithm that emulates the SDP-based approximations of nonconvex QCQPs. The cutting planes are required to be *sparse*, in order to ensure a numerically attractive approximation, and *efficiently computable*. We present a novel connection between such sparse cut generation and the sparse principal component analysis problem in statistics, which allows us to achieve these two goals. We show extensive computational results advocating for the use of our approach.

Key words. quadratically-constrained quadratic programs, nonconvex optimization, sparse cutting planes, sparse principal component analysis

AMS subject classifications. 90C26, 90C20, 90-08

1. Introduction. Nonconvex *quadratically-constrained quadratic programs* (QCQPs) are highly expressive models with wide applicability. For example, they can represent mixed-integer polynomial optimization over a compact feasible region, which already captures a broad set of real-world problems. At the same time, the flexible modeling capabilities of QCQPs imply steep computational challenges involved in designing practical general-purpose techniques for these problems.

One successful approach for solving QCQPs is based on *cutting planes*, or *cuts*, in which a more computationally tractable relaxation of the initial QCQP is formulated and then iteratively refined. Although there exist numerous families of cuts for QCQPs (see, for example, [12, 13, 17, 43, 44, 51, 55]), we focus on understanding and improving the performance of sequential *linear* cutting plane methods for solving QCQPs, which have the advantage over more complex cutting surfaces in that linear relaxations can be more easily embedded in intelligent search of the feasible region, such as spatial branch and bound [34].

One crucial aspect that has been largely underexplored for QCQPs is the effect of cut *sparsity*, in terms of the number of nonzero coefficients involved in the cut, on computational performance. Sparsity can be exploited by most modern linear programming solvers to obtain significant speedups; see for example, results and discussion in [1, 8, 25, 26, 49, 61]. While denser cuts are typically stronger, they also significantly increase the time required to solve the resulting relaxations and are associated with *tailing off* effects, in which numerical issues impair convergence.

We investigate this tradeoff between strength of dense cuts and potential speedup using sparse cuts, by developing an understanding of *which cuts* to generate and apply,

*Georgia Institute of Technology, Atlanta, GA, USA (santanu.dey@isye.gatech.edu).

†University of Florida, Gainesville, FL, USA (akazachkov@ufl.edu).

‡Canada Excellence Research Chair in Data Science for Real-Time Decision-Making. Polytechnique Montréal, Montréal, QC, Canada (andrea.lodi@polymtl.ca).

§Universidad de O'Higgins, Rancagua, Chile (gonzalo.munoz@uoh.cl).

44 with the goals of quickly improving the relaxation quality while encouraging favorable
 45 computational speed and convergence properties. Prior work has attempted to convert
 46 dense cuts into sparse ones [46]; in contrast, we draw a connection to literature from
 47 statistics on *sparse principal component analysis* (SPCA), in order to directly generate
 48 sparse cuts based on where they can best tighten the current relaxation. Sparse cuts
 49 generated in this manner have recently been shown to have promising strength in
 50 a computational study by Baltean-Lugojan et al. [5]. We substantially expand on
 51 those results through extensive computational experiments on larger-scale instances
 52 in which we evaluate sparse cuts, compared to dense ones, in terms of their strength,
 53 convergence properties, and effect on solution time.

54 Concretely, we consider nonconvex QCQPs of the form

$$55 \quad (\text{QCQP}) \quad \begin{aligned} & \min_{x \in \mathbb{R}^n} \quad x^\top Q_0 x + c_0^\top x \\ & \quad \quad \quad x^\top Q_i x + c_i^\top x + d_i \leq 0 \quad i = 1, \dots, m. \end{aligned}$$

56 The main computational difficulties that arise in solving (QCQP) are due to lack
 57 of convexity and numerical instability associated with nonlinear expressions. One
 58 method to overcome these hurdles that has garnered considerable attention is based
 59 on *semidefinite programming* (SDP).

60 The standard SDP relaxation of (QCQP), due to Shor [54], is obtained by adding
 61 new variables X_{ij} , $1 \leq i \leq j \leq n$, to replace and represent the bilinear terms $x_i x_j$ in
 62 (QCQP), and then relaxing the nonconvex condition $X = xx^\top$, yielding

$$63 \quad (1.1) \quad \begin{aligned} & \min_{x, X} \quad \langle X, Q_0 \rangle + c_0^\top x \\ & \quad \quad \quad \langle X, Q_i \rangle + c_i^\top x + d_i \leq 0 \quad i = 1, \dots, m \\ & \quad \quad \quad X - xx^\top \succeq 0. \end{aligned}$$

64 The constraint $X - xx^\top \succeq 0$ is typically expressed, by Schur's complement, as

$$65 \quad (1.2) \quad \begin{bmatrix} 1 & x^\top \\ x & X \end{bmatrix} \succeq 0.$$

66 Additionally, when some of the x variables are bounded, the SDP relaxation is then
 67 typically further refined by adding the well-known linear McCormick inequalities [41],
 68 to obtain better bounds.¹ Slightly abusing notation, we refer to the resulting relax-
 69 ation as (SDP):

$$70 \quad (\text{SDP}) \quad \begin{aligned} & \min_{x, X} \quad \langle X, Q_0 \rangle + c_0^\top x \\ & \quad \quad \quad \langle X, Q_i \rangle + c_i^\top x + d_i \leq 0 \quad i = 1, \dots, m \\ & \quad \quad \quad \text{McCormick inequalities} \\ & \quad \quad \quad \begin{bmatrix} 1 & x^\top \\ x & X \end{bmatrix} \succeq 0. \end{aligned}$$

71 It has been shown that (SDP) can provide tight approximations to the optimal
 72 value of (QCQP); some notable examples of such strong performance are the max-
 73 imum cut problem (MAXCUT) [30], optimal power flow [38], and box-constrained

¹The explicit form for the McCormick inequalities is not consequential for this paper, aside from them being linear. In fact, our results apply more generally, e.g., if the relaxation is strengthened with any valid linear inequalities for (QCQP).

74 quadratic programs (BOXQP) [2, 17, 19, 66]. We also refer the reader to [20, 62] and
 75 references therein for exactness results regarding the SDP relaxation.

76 Our goal goes beyond solving the SDP and obtaining a good dual bound: we
 77 would like to obtain an *outer approximation* of the original set that can capture the
 78 strength of the SDP relaxation, but that is lightweight and *linear*. This way, the
 79 expressiveness of an SDP relaxation could be embedded in mature branching schemes
 80 — which require relaxations that are as efficiently-solvable as possible — with ease in
 81 order to solve the original QCQP instance to provable optimality.

82 In principle, to approximate the SDP relaxation (SDP), one can exploit the ma-
 83 turity of linear programming (LP) solvers by iteratively refining LP relaxations of
 84 (SDP) in a cutting-plane fashion [31, 48, 53]. Consider a first LP relaxation of (SDP),
 85 obtained by removing the SDP constraint:

$$\begin{aligned} & \min_{x, X} \quad \langle X, Q_0 \rangle + c_0^\top x \\ 86 \quad (\text{LP}) \quad & \langle X, Q_i \rangle + c_i^\top x + d_i \leq 0 \quad i = 1, \dots, m \\ & \text{McCormick inequalities.} \end{aligned}$$

87 Let (\tilde{x}, \tilde{X}) denote an optimal solution to this relaxation. Throughout the paper, we
 88 will refer to the matrix $M(x, X)$ for a given $(x, X) \in \mathbb{R}^n \times \mathbb{R}^{\binom{n}{2}+n}$, which we define
 89 as

$$90 \quad (M) \quad M(x, X) := \begin{bmatrix} 1 & x^\top \\ x & X \end{bmatrix}.$$

91 For notational convenience, we define

$$92 \quad (\tilde{M}) \quad \tilde{M} := M(\tilde{x}, \tilde{X}).$$

93 If $\tilde{M} \succeq 0$, then (SDP) is solved. Otherwise, we can efficiently find a vector
 94 $v \in \mathbb{R}^{n+1}$ such that

$$95 \quad v^\top \tilde{M} v < 0,$$

96 e.g., an eigenvector of \tilde{M} with negative eigenvalue. Then, the *linear* inequality

$$97 \quad (1.3) \quad \left\langle vv^\top, \begin{bmatrix} 1 & x^\top \\ x & X \end{bmatrix} \right\rangle \geq 0$$

98 is valid for (SDP) and cuts off (\tilde{x}, \tilde{X}) . This procedure can be viewed as a finite
 99 approximation of the semi-infinite outer description of the *positive semidefinite* (PSD)
 100 cone:

$$101 \quad (1.4) \quad M \succeq 0 \iff \langle vv^\top, M \rangle \geq 0 \quad \forall v \in \mathbb{R}^{n+1}.$$

102 This family of cutting planes has been considered before (see [5, 6, 31, 46, 48, 53]),
 103 and the main drawback has been repeatedly acknowledged: the vector v will typically
 104 be *dense*, which can cause numerical instability in the LP after adding many cuts of
 105 the form (1.3). Fukuda et al. [29] provide one way of avoiding this issue, in the presence
 106 of *structured sparsity* of the objective and linear constraints in (SDP). Under such
 107 structured sparsity, one can (without any loss in relaxation quality) replace (1.2), the
 108 PSD requirement on the full matrix $M(x, X)$, with the requirement that a set of *small*
 109 principal submatrices of $M(x, X)$ is PSD, which implies that it suffices to only add
 110 the cuts of type (1.3) associated with the eigenvectors of those principal submatrices.

111 We instead follow a different direction, in which we directly enforce a target
 112 sparsity, say k , on the vector v . This involves searching for a v such that

$$113 \quad (1.5) \quad v^\top \tilde{M} v < 0 \quad \text{and} \quad \|v\|_0 := |\text{supp}(v)| \leq k,$$

114 where $\text{supp}(v)$ denotes the nonzero entries of v and $|\cdot|$ is the cardinality operator. We
 115 call a vector $v \in \mathbb{R}^{n+1}$ a k -sparse-eigencut if it satisfies (1.5), and if the k -length
 116 vector obtained by only taking the nonzero elements of v is a unit eigenvector of the
 117 principal submatrix of \tilde{M} defined on the indices in $\text{supp}(v)$.

118 *Contributions.* We formulate the problem of finding a k -sparse-eigencut, or de-
 119 termining if none exists, as an optimization problem and show that it is equivalent to
 120 Sparse Principal Component Analysis (SPCA). This implies the problem is \mathcal{NP} -hard,
 121 and from the results of Blekherman et al. [9], we observe that there exist matrices
 122 where k -sparse-eigencuts exist, but there are no $(k-1)$ -sparse-eigencuts. In spite of
 123 these negative worst-case results, the same connection with SPCA allows us to use
 124 an empirically-strong heuristic for this problem to efficiently compute one strong k -
 125 sparse-eigencut in practice. We then devise a novel scheme to compute multiple k -
 126 sparse-eigencuts for a given \tilde{M} , and we conduct extensive computational experiments
 127 using a cutting plane approach. Our results strongly advocate for the importance of
 128 sparsity, and show that a lightweight polyhedral approach can successfully approxi-
 129 mate the quality of SDP-based approaches.

130 *Notation.* We define $[n] := \{1, \dots, n\}$. The cardinality of a set I is denoted by
 131 $|I|$. We denote the set of n -dimensional vectors with real entries by \mathbb{R}^n , and the set
 132 of $n \times n$ real-valued matrices by $\mathbb{R}^{n \times n}$. For $v \in \mathbb{R}^n$, $\|v\|_0$ denotes the number of
 133 nonzero entries in v . We use $\langle \cdot, \cdot \rangle$ to represent the matrix inner-product. For a matrix
 134 $X \in \mathbb{R}^{n \times n}$ and $I \subseteq [n]$, we let X_I be the principal submatrix of X given by the
 135 columns and rows indexed by I . Similarly, v_I will be the vector corresponding to the
 136 entries of v indexed by I . We denote by \mathcal{S}_+^n the cone of $n \times n$ PSD matrices. For an
 137 $n \times n$ matrix X , we also use $X \succeq 0$ to denote $X \in \mathcal{S}_+^n$. We let $\mathcal{S}_+^{n,k}$ denote the cone
 138 of $n \times n$ matrices such that every $k \times k$ principal submatrix is PSD; that is,

$$139 \quad \mathcal{S}_+^{n,k} := \{X \in \mathbb{R}^{n \times n} : X_I \in \mathcal{S}_+^k, \forall I \subseteq [n], |I| = k\}.$$

140 **2. Literature review.** Given the vast literature involving cutting planes for
 141 nonconvex QCQPs, here we restrict ourselves to reviewing approaches which rely on
 142 the structure

$$143 \quad (2.1) \quad X = xx^\top, x \in [\ell, u],$$

144 in order to derive inequalities. We refer the reader to [12] for a survey on cutting
 145 planes methods for nonlinear optimization.

146 Historically, the main limitation of using the SDP relaxation (SDP) has been
 147 its lack of scalability, but significant progress has been made on this front; see, for
 148 example, [18, 68]. One way to alleviate the computational burden of the SDP is
 149 leveraging *structured sparsity*. When the bilinear terms are represented as edges of a
 150 graph, and this graph is close to being a tree, one can use the framework of Fukuda
 151 et al. [29] to avoid imposing (1.2) on a large matrix, and instead enforce the PSD
 152 requirement over small submatrices. We refer the reader to [36, 37, 60] and references
 153 therein for other approaches exposing structured sparsity from an SDP perspective.

154 The most common way to exploit (2.1) is by adding the McCormick inequalities
 155 [41]. These are valid inequalities derived from each equality $X_{i,j} = x_i x_j$ and variable

156 bounds. Using these inequalities can provide strong relaxations in combination with
 157 the SDP constraint (1.2) [2, 16, 17, 19, 66].

158 More generally, in the presence of (2.1), one can use valid inequalities for the
 159 Boolean Quadric Polytope [45] in order to obtain valid inequalities for (QCQP). Their
 160 strong computational performance for BOXQP was shown by Bonami et al. [11]. See
 161 also [14, 17, 65] and [13, 23, 27, 47] for cuts that are valid for the convex hull of side
 162 constraints together with (2.1).

163 An alternative way to generate cuts in the presence of (2.1) is through the con-
 164 struction of convex sets whose interior *does not* intersect the feasible region. These
 165 sets can be used via the intersection cut [4, 58] framework in order to construct valid
 166 inequalities. The work in the papers [6, 7, 22, 44] falls in this category. Bienstock
 167 et al. [7] introduce the concept of *outer-product-free sets* as convex sets of matrices
 168 that do not include any matrix of the form $X = xx^\top$ in its interior. Fischetti and
 169 Monaci [28] construct *bilinear-free* sets through a bound disjunction and McCormick
 170 inequalities.

171 Beyond linear programming relaxations, Saxena et al. [52] use eigenvectors of
 172 $X - xx^\top$ to construct *convex quadratic* inequalities that iteratively approximate the
 173 constraint $X - xx^\top = 0$.

174 In contrast to the techniques mentioned above, which would cut through the
 175 constraint (1.2), we derive valid inequalities for the semidefinite *relaxation* of $X = xx^\top$
 176 in (2.1), in which this constraint is replaced with (1.2). The inequalities we construct
 177 are valid for $\{(x, X) : M(x, X) \succeq 0\}$, and additionally are required to be sparse. There
 178 are two papers that propose a similar methodology. Qualizza et al. [46] look for k -
 179 sparse-eigencuts by iteratively setting to zero the components of a dense eigenvector
 180 until reaching the target sparsity. Baltean-Lugojan et al. [5] use a neural network in
 181 order to select a promising $k \times k$ submatrix of $M(x, X)$, for $k \leq 6$, to compute a k -
 182 sparse-eigencut for the problem. A theoretical analysis on the quality achieved by
 183 k -sparse-eigencuts was recently provided in [9, 10]. The authors provide upper and
 184 lower bounds on the distance between \mathcal{S}_+^n and $\mathcal{S}_+^{n,k}$. Recently, Rodrigues de Sousa
 185 et al. [50] used eigenvector-based cuts for obtaining strong linear relaxations to the
 186 maximum k -cut problem.

187 **3. The k -sparse separation problem and SPCA.** We seek a k -sparse-eigencut
 188 violated by (\tilde{x}, \tilde{X}) , which can be thought of as an optimal solution to (LP) or a dif-
 189 ferent relaxation, in which $\tilde{M} \not\succeq 0$. This is equivalent to determining if $\tilde{M} \in \mathcal{S}_+^{n,k}$,
 190 since it can easily seen that

$$191 \quad A \in \mathcal{S}_+^{n,k} \iff v^\top Av \geq 0 \quad \forall v \in \mathbb{R}^n, \|v\|_0 \leq k.$$

192 Finding a k -sparse-eigencut, or determining if none exists, is closely related to
 193 the SPCA problem, which we define next.

194 **DEFINITION 3.1.** *Given a matrix $A \in \mathcal{S}_+^{n+1}$ and a sparsity level $k \in \mathbb{N}$, the k -*
 195 *Sparse Principal Component Analysis (k -SPCA) problem is defined as*

$$196 \quad (3.1) \quad \begin{aligned} & \max_{v \in \mathbb{R}^{n+1}} v^\top Av \\ & \|v\|_2 = 1 \\ & \|v\|_0 \leq k. \end{aligned}$$

198 *Its decision version reads: given $A \in \mathcal{S}_+^{n+1}$, $k \in \mathbb{N}$, and $K \in \mathbb{R}_{\geq 0}$, determine if*

199 there exists $v \in \mathbb{R}^n$ such that

$$200 \quad (3.2) \quad \|v\|_2 = 1, \quad \|v\|_0 \leq k, \quad v^\top A v > K.$$

201 We begin by showing that one can reduce the decision version of the SPCA
202 problem to determining if a k -sparse-eigencut exists. The proof is trivial.

203 PROPOSITION 3.2. Let $A \in \mathcal{S}_+^{n+1}$, $k \in \mathbb{N}$ and $K \in \mathbb{R}$. A vector v satisfies (3.2) if
204 and only if v is a k -sparse-eigencut of \tilde{A} defined as

$$205 \quad \tilde{A} := KI - A,$$

206 where I is the $(n+1) \times (n+1)$ identity.

207 This immediately implies that finding a k -sparse-eigencut is \mathcal{NP} -hard [40, 56, 57].
208 Although we cannot hope for an efficient algorithm for finding our proposed cuts,
209 SPCA is a well-studied problem and efficient practical heuristics exist [3, 24, 35, 39,
210 67]. We pursue the connection of our separation problem and SPCA through the
211 problem of finding the *most violated* k -sparse-eigencut:

$$212 \quad (3.3) \quad \begin{aligned} & \min_{v \in \mathbb{R}^{n+1}} v^\top \tilde{M} v \\ & \|v\|_2 = 1 \\ & \|v\|_0 \leq k. \end{aligned}$$

213
214 If the value of (3.3) is negative, we obtain a valid k -sparse-eigencut, and otherwise,
215 none exists. Whenever the matrix \tilde{M} is negative semidefinite, problem (3.3) is exactly
216 a k -SPCA problem as (3.1). In our case, however, we are interested in studying the
217 problem with no assumption over \tilde{M} . Lemma 3.3 shows that, in any case, (3.3) is
218 equivalent to (3.1).

219 LEMMA 3.3. For every $(\tilde{x}, \tilde{X}) \in \mathbb{R}^n \times \mathbb{R}^{\binom{n}{2}+n}$ and $k \in \mathbb{N}$, there exists a matrix
220 $A \succeq 0$ such that (3.3) is equivalent to solving the k -SPCA problem (3.1).

221 *Proof.* Let λ^{\max} be the largest eigenvalue of \tilde{M} . We can equivalently rewrite (3.3)
222 as

$$223 \quad \begin{aligned} & \min_v v^\top (\tilde{M} - \lambda^{\max} I) v + \lambda^{\max} v^\top v \\ & \|v\|_2 = 1 \\ & \|v\|_0 \leq k, \end{aligned}$$

227 with I the $(n+1) \times (n+1)$ identity matrix. Since $\|v\|_2 = 1$, the above is equivalent
228 to

$$229 \quad \begin{aligned} & \lambda^{\max} - \max_v v^\top (\lambda^{\max} I - \tilde{M}) v \\ & \|v\|_2 = 1 \\ & \|v\|_0 \leq k. \end{aligned}$$

233 It easy to see that $\lambda^{\max} I - \tilde{M}$ is positive semidefinite. We define $A = \lambda^{\max} I - \tilde{M}$ and
234 conclude that (3.3) is equivalent to an SPCA problem. \square

235 Due to the \mathcal{NP} -hardness results discussed above, it is unlikely there exists a
236 polynomial-time algorithm to compute a separating k -sparse-eigencut. However, one
237 could still hope to always find a good cut independently of the computing time in-
238 volved. The following example, constructed by Blekherman et al. [9], shows that this
239 cannot be done in general.

240 EXAMPLE 3.4. Consider the $n \times n$ matrix $G(a, b)$ whose entries are b in the di-
 241 agonal entries and $-a$ in the off-diagonal entries, for some $a, b \geq 0$. That is, $G(a, b)$
 242 takes the form

$$243 \quad G(a, b) = \begin{bmatrix} b & -a & \cdots & -a \\ -a & b & \cdots & -a \\ \vdots & & \ddots & \vdots \\ -a & \cdots & -a & b \end{bmatrix}.$$

244 The results in [9] show that if a, b are such that

$$245 \quad (k-1)a \leq b < (n-1)a,$$

246 then all $k \times k$ submatrices of $G(a, b)$ are PSD, but $G(a, b) \not\succeq 0$. In other words, even if
 247 a given $n \times n$ matrix is not PSD, there is no guarantee that an $(n-1)$ -sparse-eigencut
 248 exists. ■

249 In Example 3.4, the matrix $G(a, b)$ is dense, and one may naturally wonder if the
 250 same result holds for a sparse matrix, since in this case one might expect it would be
 251 easier to find sparse cuts. The next example shows that no such hope can be realized
 252 in general.

253 EXAMPLE 3.5. Consider now a $n \times n$ matrix defined as $G(a, b)$ and let $m \in \mathbb{N}$.
 254 Let $N = nm$ and define the block diagonal $N \times N$ matrix $\tilde{G}(a, b)$ whose blocks are
 255 given by $G(a, b)$, with sparsity $1/m$.

256 Since $\tilde{G}(a, b) \succeq 0$ if and only if $G(a, b) \succeq 0$, it suffices to consider

$$257 \quad (n-2)a \leq b < (n-1)a$$

258 in order to obtain that $\tilde{G}(a, b) \not\succeq 0$. In this case, the existence of n -sparse-eigencuts
 259 is guaranteed, but there is no $(n-1)$ -sparse-eigencut. ■

260 While the above examples are negative results, they are only informing us that
 261 we cannot hope to devise an efficient method that will work in the worst case. In
 262 practice, the story might be different. In the following, we first analyze the empirical
 263 expressiveness of these cuts and then exploit the connection with SPCA in an efficient
 264 *on-the-fly* generation of sparse cuts. Specifically, in Section 4 we study how much
 265 we can expect dual bounds to improve if we have access to all k -sparse-eigencuts;
 266 this motivates our main algorithm (Section 5), which computes k -sparse-eigencuts
 267 sequentially. These computations, due to the connection drawn by Lemma 3.3, can
 268 be transformed into solving sequences of SPCA instances, for which we can exploit
 269 efficient heuristics from the SPCA literature.

270 4. Empirical expressiveness of k -sparse-eigencuts. The computational hard-

271 ness of solving (3.3) to find a k -sparse-eigencut lies in selecting an optimal *support*, i.e.,
 272 where the nonzeros in the solution should be. If the support of size k of an optimal
 273 solution of (3.3) is known, finding the actual optimal solution reduces to computing
 274 an eigendecomposition of the corresponding $k \times k$ submatrix of \tilde{M} . This support
 275 selection is what Baltean-Lugojan et al. [5] tackle via a neural network. Their results
 276 suggest that k -sparse-eigencuts can substantially close the gap between the optimal
 277 values of (LP) and the SDP relaxation with McCormick inequalities added. However,
 278 for our purposes, we desire a more refined analysis of the *expressiveness*, or modeling
 279 power, of k -sparse-eigencuts, particularly in comparison to the performance of *dense*
 280 eigencuts.

281 In this section, we perform experiments with low-dimensional instances for which
 282 we can exhaustively enumerate all $\binom{n+1}{k}$ possible supports, to compute cuts via the
 283 eigendecompositions of all possible $k \times k$ principal submatrices of \tilde{M} . Our goal is to
 284 answer the following questions:

- 285 *Question 1.* What is the right number of k -sparse-eigencuts to add?
 286 *Question 2.* What is the appropriate level of sparsity, k , to use?
 287 *Question 3.* Given a budget on the number of cuts we can add, can we efficiently
 288 identify a set of strong k -sparse-eigencuts?

289 To this end, first, in Section 4.3.1, we evaluate how much a *single* k -sparse-eigencut
 290 can improve the dual bound in an optimization problem. Then, in Section 4.3.2, we
 291 assess how much the dual bound improves if we add multiple k -sparse-eigencuts,
 292 including *all* possible cuts of this type, or only a limited number of cuts. To answer
 293 Question 3, we examine metrics by which a good set of eigencuts can be selected,
 294 in order to find a few k -sparse-eigencuts that are responsible for most of the bound
 295 improvement.

296 The answer to these questions motivate the choices we make in our main algo-
 297 rithm, Algorithm 1, which is presented in Section 6. Specifically, the enumeration
 298 experiments justify how many sparse cuts we should generate at each round, and
 299 whether or not we should allow the use of a few dense eigencuts.

300 **4.1. Experimental setup for enumeration experiments.** We consider three
 301 low-dimensional instances of the BOXQP library [15, 59]: `spar30-060-1`, `spar30-080-1`
 302 and `spar030-100-1`; these have the form

$$303 \quad \min_x x^\top Qx + c^\top x$$

$$304 \quad x \in [0, 1]^n$$

306 with $Q \not\prec 0$. For these instances, $n = 30$, and the second number in their names (60,
 307 80, and 100) indicates (roughly) the percentage of nonzero entries in Q . While for
 308 succinctness we report on only these three instances, we obtained similar outcomes
 309 with other instances from the family.

310 For each instance, we first solve the LP relaxation (LP) of (SDP). As before,
 311 let (\tilde{x}, \tilde{X}) be an optimal solution to this LP. Then, we enumerate all possible $\binom{n+1}{k}$
 312 supports for $k = 4$, and compute eigenvectors for all negative eigenvalues of the
 313 corresponding $k \times k$ submatrix of \tilde{M} . We also compute *dense* cuts, which are just
 314 eigenvectors of all negative eigenvalues of \tilde{M} . We only add one round of cuts. A *round*
 315 of cuts is obtained from a given LP solution, without adding any new cuts to the LP
 316 and reoptimizing to get a new solution.

317 We implemented the enumeration code in C++. We use the Eigen library [32] for
 318 linear algebra computations and the LP solver is Gurobi 9.0 [33]. In order to measure
 319 relaxation quality, we also solve the SDP relaxation (SDP) using the C++ Fusion API
 320 of MOSEK version 9.2 [42].

321 **4.2. Performance measures.** We measure strength through the commonly
 322 used *gap closed*. Let us denote the initial (LP) optimal value as LP_{opt} and the SDP
 323 optimal value of (SDP) as SDP_{opt} , which is an upper bound on the value any set
 324 of eigencuts can achieve. For a subsequent LP relaxation LP' , obtained by adding
 325 additional cuts to the base LP, the gap closed $GC(LP')$ is

$$326 \quad GC(LP') = 100 \times \frac{\text{LP}'_{\text{opt}} - \text{LP}_{\text{opt}}}{\text{SDP}_{\text{opt}} - \text{LP}_{\text{opt}}}.$$

TABLE 1

Percent gap closed by k -sparse-eigencuts from a single support, contrasted with requiring a single $k \times k$ principal submatrix to be PSD, and with adding dense cuts.

Instance	1-supp-cuts		1-supp-PSD		Dense cuts
	Max	Avg	Max	Avg	
spar30-060-1	3.37	0.25	3.69	0.31	38.77
spar30-080-1	4.19	0.57	4.83	0.75	49.31
spar30-100-1	4.56	0.95	4.95	1.20	60.94

327 Given an optimal solution (\tilde{x}, \tilde{X}) to the initial LP relaxation and an eigencut
 328 of the form (1.3), that is, $\langle vv^\top, \tilde{M} \rangle \geq 0$, the *violation* is the nonnegative number
 329 $-\langle vv^\top, \tilde{M} \rangle$. In all our experiments, $\|v\|_2 = 1$, making the violation of different cutting
 330 planes comparable.

331 While we focus on relaxation strength in this section, we apply more comprehen-
 332 sive quality metrics, including solution time, in the more extensive experiments of
 333 Section 6.

334 4.3. Results.

335 **4.3.1. Single support reports.** In Table 1, for $k = 4$, we summarize how much
 336 gap is closed by three approaches: adding k -sparse-eigencuts obtained from a single
 337 $k \times k$ principal submatrix, requiring that a $k \times k$ principal submatrix is PSD, and
 338 adding dense cuts. The columns labeled *1-supp-cuts* show the maximum and average
 339 gap closed when all k -sparse-eigencuts for a single $k \times k$ support are added to (LP).
 340 The columns labeled *1-supp-PSD* show the maximum and average gap closed when
 341 a single $k \times k$ support is imposed to be PSD. The column labeled *Dense cuts* shows
 342 the gap closed by adding all dense cuts.

343 From this table, we draw several conclusions. First, we see that convexifying a
 344 single $k \times k$ principal submatrix has only limited impact on gap closed. Adding a
 345 PSD requirement for one $k \times k$ submatrix never closes more than 5% of the gap, and
 346 the average is usually around 1%. The corresponding k -sparse-eigencuts also do not
 347 perform well in this case, though the gap closed from a round of cuts is not significantly
 348 different from the gap closed by imposing a single submatrix to be PSD. Second, the
 349 performance of dense cuts is remarkable: a large percentage of the gap is closed with
 350 only one round of cuts, which for these instances is 14–16 cuts. Lastly, there is a
 351 trend that, overall, cuts are more effective in dense instances. This is somewhat
 352 expected: the larger the number of zeros in the objective, the more common it is for
 353 $k \times k$ submatrices to have a considerable portion of zero objective coefficients. In this
 354 case, if a sparse support is convexified, the objective value may not change by much.
 355 However, the strength of sparse cuts from a single support, as a proportion of the gap
 356 closed by dense cuts, is higher in sparser instances.

357 **4.3.2. Multiple support reports.** The single-support experiments suggest that, ■
 358 in order to have some impact with k -sparse-eigencuts, adding cuts across many sup-
 359 ports simultaneously is necessary. This raises the question of which supports to use,
 360 given that, in practice, we have a budget on the number of cuts we can add. We
 361 examine this in our next set of experiments, in which we evaluate different support
 362 selection criteria. We assign a score to each of the 31,465 supports, then select the
 363 top 5% of supports (1,574 supports) having the highest score.

TABLE 2

Percent gap closed when multiple supports are used to generate k -sparse-eigencuts, contrasted with requiring the corresponding principal submatrices to be PSD, and with adding dense cuts.

Instance	k -cuts			k -PSD		Dense cuts
	All	Top (gap)	Top (viol)	All	Top (psd)	
spar30-060-1	50.51	31.46	18.93	61.31	44.38	38.77
spar30-080-1	80.95	63.20	42.03	91.53	84.08	49.31
spar30-100-1	87.84	63.29	66.01	95.32	82.51	60.94

364 The scores we test for each support are: (1) *gap*, which is the change in objective
365 value when all k -sparse-eigencuts from that support are added; (2) *violation*, which
366 is the maximum violation of any k -sparse-eigencut from that support; and (3) *psd*,
367 which denotes the change in objective value after adding the PSD requirement on that
368 support. We include both gap and violation as scores, because violation is commonly
369 used to measure the importance of a cut, due to it being cheap to compute relative
370 to resolving an LP; hence, we are interested in evaluating whether violation is a good
371 proxy for expected objective improvement in this setting.

372 Table 2 shows the results. Columns 2–4, with the heading k -cuts, give the gap
373 closed by k -sparse-eigencuts when all possible such cuts are added, and when the top
374 5% of cuts are added, sorted by “gap” and “violation”. Columns 5 and 6, with the
375 heading k -PSD, show the gap closed when requiring $k \times k$ principal submatrices are
376 PSD, both for all such matrices, and for only the top 5% sorted by their “psd” score.
377 Column 7 repeats the gap closed by dense cuts from Table 1.

378 From Table 2, we observe many interesting phenomena. First, we see that adding
379 all sparse cuts significantly outperforms all dense cuts; however, while only 14–16
380 dense cuts are added per instance, the number of sparse cuts is in the thousands. On
381 one hand this indicates sparse cuts can replicate the strength of dense cuts, but on
382 the other hand, any speed advantages from resolving the LP with sparser cuts are
383 likely to be negated by that many additional constraints. Second, adding cuts from
384 only the top 5% of supports can achieve 60 to 77% of the gap closed by all supports.
385 This indicates that, although multiple supports need to be considered to produce a
386 strong relaxation, an intelligent choice of a small number of supports can suffice. Last,
387 concerning the violation measure: while, on average, the selection of cuts via violation
388 performs worse than its gap closed counterpart, it can happen, as in `spar30-100-1`,
389 that the top 5% of supports sorted by violation together close more gap than if the
390 top 5% of supports were chosen by their individual objective improvement.

391 **4.4. Summary of enumeration experiments.** We conclude this section with
392 the high-level messages imparted by these experiments. First, k -sparse-eigencuts
393 closely emulate, in strength, the addition of the PSD requirement to a support. Sec-
394 ond, no single support has a significant effect in terms of gap closed. Therefore, one
395 should find multiple supports that together perform well. Third, the violation of a cut
396 presents a good and efficient heuristic for cut selection in this setting. Finally, dense
397 cuts are remarkably strong, considering there are usually very few of them available.

398 We do not directly consider interactions among supports: that is, we score in-
399 dividual k -length subsets of $[n + 1]$, whereas it may be better to score, for example,
400 pairs of supports, or otherwise incorporating combinatorial effects.

401 Since the experiments in this section required enumeration of all $\binom{n+1}{k} \in \Theta((n +$

402 $1)^k$) possible sparse supports, the technique is not practical for even medium-size
 403 instances or moderate values of k . The same drawback is observed by Baltean-Lugojan
 404 et al. [5]; their neural-network-based approach can estimate if a *given* support will
 405 have a considerable effect (in terms of the cut generated), but they still need to
 406 enumerate supports for their evaluation. Our next goal is to effectively choose strong
 407 sets of sparse supports *on the fly*.

408 **5. Computation of multiple sparse cuts.** We develop an algorithm to com-
 409 pute multiple k -sparse-eigencuts for a given matrix \tilde{M} , using an oracle that can com-
 410 pute a single k -sparse-eigencut. We compare the performance of the cutting planes
 411 produced by this algorithm with the results of the previous low-dimensional enumer-
 412 ation experiments.

413 **5.1. Iterative k -sparse-eigencut computation.** Our strategy to compute
 414 multiple sparse cuts starts with an oracle that efficiently returns a single cut with
 415 sparsity level k . Instead of using that cut directly, we take its support, say I , and we
 416 add all the k -sparse-eigencuts from \tilde{M}_I .² Lemma 5.1 formally states that such cuts
 417 exist from \tilde{M}_I .

418 **LEMMA 5.1.** *Let $\tilde{M} \in \mathbb{R}^{(n+1) \times (n+1)}$ be a symmetric matrix, and suppose $w \in$
 419 \mathbb{R}^{n+1} is such that $w^\top \tilde{M} w < 0$. Let $I := \text{supp}(w)$ and $k := |I|$. The number of k -
 420 sparse-eigencuts of \tilde{M}_I is between 1 and k .*

421 *Proof.* The number of k -sparse-eigencuts from \tilde{M}_I depends on its number of neg-
 422 ative eigenvalues. The result follows because the matrix \tilde{M}_I has at most k dis-
 423 tinct unit eigenvectors, and at least one of those is associated with a negative ei-
 424 genvalue. This second fact follows from the existence of a sparse cut using w , since
 425 $0 > w^\top \tilde{M} w = w_I^\top \tilde{M}_I w_I$, so that \tilde{M}_I is not positive semidefinite. \square

426 Based on the results in Section 4, we know that limiting ourselves to one support
 427 may lead to a weak set of cuts. The next question is how to use \tilde{M} to generate more
 428 cuts, having exhausted the ones from the initial support I . A natural first idea is to
 429 target a set of indices from $[n+1]$ that is disjoint to I , to diversify the areas of \tilde{M} being
 430 convexified. We implemented a version of this proposal, and on small- to medium-size
 431 BOXQP instances, there were encouraging results. However, this approach is unable
 432 to adequately capitalize on another observation from Section 4, that a sufficiently
 433 large number of sparse cuts can outperform dense cuts. For example, for a moderate
 434 sparsity level such as $k = n/4$, the orthogonal scheme can only generate up to 4
 435 supports and n cuts (by Lemma 5.1), while we desire more cuts from each iteration.

436 With that motivation, we present Algorithm 1, our strategy to compute multi-
 437 ple sparse cuts from \tilde{M} , without solving another LP. The oracle referenced within
 438 Algorithm 1 solves (3.3). It can be helpful for intuition to keep in mind the case
 439 that $k = n + 1$, for which Algorithm 1 simply returns all eigenvectors with a negative
 440 eigenvalue in an eigendecomposition of the matrix.

441 We now establish finiteness of this algorithm.

442 **LEMMA 5.2.** *Algorithm 1 terminates in a finite number of iterations, even if the*
 443 *parameter MAXNUMSUPPORTS is set to infinity.*

444 *Proof.* At each start of the while loop, the matrix M_I^i has signature³ (p, q) with
 445 $q \geq 1$ (otherwise, the algorithm would have terminated), and M_I^{i+1} has signature

²Adding all k -sparse-eigencuts, as opposed to just one per support, is shown to be effective in the computational results of Qualizza et al. [46] with their *Minor PSD cuts*.

³The signature of a matrix is (p, q) if it has p positive and q negative eigenvalues.

Algorithm 1: SPARSEROUND(\tilde{M}, k): one round of k -sparse-eigencuts

Input : A matrix $\tilde{M} \in \mathbb{R}^{(n+1) \times (n+1)}$ with $\tilde{M} \not\geq 0$, and a sparsity level $k \in \mathbb{N}$.
Parameters: MAXNUMSUPPORTS: maximum number of considered supports.
ORACLE: oracle for solving (3.3).
Output : A sequence of k -sparse-eigencuts $\{\tilde{w}_j\}_{j=1}^p$ such that $\|\tilde{w}_j\|_2 = 1$, $\|\tilde{w}_j\|_0 \leq k$,
and $\tilde{w}_j^\top \tilde{M} \tilde{w}_j < 0$, for $j \in [p]$.

- 1 **Initialize:** $p \leftarrow 0$, $i \leftarrow 1$, $M^1 \leftarrow \tilde{M}$, and $w \leftarrow \text{ORACLE}(M^1)$;
- 2 **while** $w^\top M^i w < 0$ and $i < \text{MAXNUMSUPPORTS}$ **do**
- 3 $I \leftarrow \text{supp}(w)$;
- 4 Let λ_i^{\min} and q_i denote the most negative eigenvalue, and associated unit eigenvector,
of M_I^i ;
- 5 Let \tilde{w}_i denote q_i lifted to \mathbb{R}^{n+1} by setting all components not in I to 0;
- 6 $M^{i+1} \leftarrow M^i - \lambda_i^{\min} \tilde{w}_i \tilde{w}_i^\top$;
- 7 $i \leftarrow i + 1$ and $p \leftarrow p + 1$;
- 8 $w \leftarrow \text{ORACLE}(M^i)$;
- 9 **end**
- 10 **return** $\{\tilde{w}_j\}_{j=1}^p$;

446 $(p, q-1)$. Additionally, if an arbitrary $k \times k$ principal submatrix M_S^i , with $S \subseteq [n+1]$,
447 $|S| = k$, has signature (p, q) , then M_S^{i+1} has signature (p', q') with $q' \leq q$. This is
448 because M_S^{i+1} is obtained from adding a PSD matrix $(-\lambda_i^{\min} \tilde{w}_S \tilde{w}_S^\top)$, with \tilde{w}_S the
449 sub-vector of \tilde{w}_i given by the entries in S) to M_S^i .

450 Thus, at each step, the number of negative eigenvalues of *every* $k \times k$ submatrix
451 does not increase, and decreases strictly for at least one such submatrix. Since there
452 are at most $\binom{n+1}{k}$ principal $k \times k$ submatrices, the algorithm finishes in at most $k \binom{n+1}{k}$
453 steps. \square

454 The following lemma shows that all k -sparse-eigencuts generated by Algorithm
455 1 are valid inequalities being violated by \tilde{M} . Additionally, it shows precisely which
456 matrices M^i are being cut by each generated k -sparse-eigencut.

457 **LEMMA 5.3.** *The sequence of vectors $\{\tilde{w}_j\}_{j=1}^p$ generated by Algorithm 1 satisfies*

- 458 1. $\tilde{w}_i^\top M^i \tilde{w}_i < 0$ for every $i \in [p]$,
- 459 2. $\tilde{w}_i^\top \tilde{M} \tilde{w}_i < 0$ for every $i \in [p]$, and
- 460 3. $\tilde{w}_i^\top M^{i+1} \tilde{w}_i = 0$ and $\tilde{w}_j^\top M^{i+1} \tilde{w}_j > 0$, for $1 \leq j \leq i-1 \leq p$.

461 *Proof.* Part 1 follows by assumption of the existence of a cut returned by ORACLE.
462 Part 2 follows by noting that

$$463 \quad \tilde{M} = M^i + \sum_{j=1}^{i-1} \lambda_j^{\min} \tilde{w}_j \tilde{w}_j^\top.$$

464 Therefore,

$$465 \quad \tilde{w}_i^\top \tilde{M} \tilde{w}_i = \tilde{w}_i^\top M^i \tilde{w}_i + \sum_{j=1}^{i-1} \lambda_j^{\min} (\tilde{w}_i^\top \tilde{w}_j)^2 < 0,$$

466 where the last inequality follows from Part 1 and the fact that $\lambda_j^{\min} < 0$. The first
467 statement in Part 3 follows since $\|\tilde{w}_i\|_2 = 1$ and it is the eigenvector associated to
468 λ_i^{\min} , thus

$$469 \quad \tilde{w}_i^\top M^{i+1} \tilde{w}_i = \tilde{w}_i^\top M^i \tilde{w}_i - \lambda_i^{\min} = 0.$$

TABLE 3

Summary of gap closed by multiple k -sparse-eigencuts selected by Algorithm 1, contrasted with same number of cuts selected via enumeration and sorting and imposing PSD-ness.

Instance	k -cuts		k -PSD
	Algorithm 1	Top (gap)	Top (psd)
spar30-060-1	5.48	9.06	10.01
spar30-080-1	14.98	16.83	25.60
spar30-100-1	22.62	25.48	30.52

470 For the second statement in Part 3, we proceed via induction over $i - j \geq 1$. For
471 $j = i - 1$,

$$472 \quad \tilde{w}_{i-1}^\top M^{i+1} \tilde{w}_{i-1} = \tilde{w}_{i-1}^\top M^i \tilde{w}_{i-1} - \lambda_i^{\min} (\tilde{w}_i^\top \tilde{w}_{i-1})^2 = -\lambda_i^{\min} (\tilde{w}_i^\top \tilde{w}_{i-1})^2 > 0.$$

473 Lastly, for general $i - j$,

$$474 \quad \tilde{w}_j^\top M^{i+1} \tilde{w}_j = \tilde{w}_j^\top M^i \tilde{w}_j - \lambda_i^{\min} (\tilde{w}_i^\top \tilde{w}_j)^2 > 0,$$

475 where the inequality follows by the inductive step. \square

476 **5.2. Enumeration experiments revisited.** To get a sense of the quality of
477 the cuts produced by Algorithm 1, we compare them to the other selection procedures
478 evaluated in Section 4.3.2.

479 Our implementation of ORACLE used in Algorithm 1 is based on the *Truncated*
480 *Power Method* (TPower) by Yuan and Zhang [67], which is an efficient practical
481 heuristic, with some theoretical guarantees, to generate high-quality solutions for the
482 SPCA problem. Specifically, we run TPower after appropriately modifying the current
483 solution \tilde{M} as per our discussion in Lemma 3.3.

484 The results of the experiments are shown in Table 3. For these small examples,
485 Algorithm 1 does not generate as many cuts as generated by the scores we considered
486 in the earlier enumeration experiments: the TPower method may fail to yield a cut
487 even if one is available, and step 6 of Algorithm 1 may inadvertently discard some
488 supports from which a cut could be computed. Note that the gaps closed in Table 2
489 are considerably larger than those of Table 3, but the latter requires a full enumeration
490 of supports. For this reason, we restrict the number of supports considered for the
491 “top” selection in Table 3 to be the same number of supports used by Algorithm 1.

492 We see that Algorithm 1 performs quite well. It is always within 1% of the
493 gap closed by the supports selected by “violation”, and less than 4% away from the
494 gap-closed-based selection. Unlike the other scores, Algorithm 1 does not require a
495 complete enumeration of the supports, and it does not require us to solve an LP or
496 SDP: it achieves our goal of generating supports dynamically.

497 At this point, we have motivated the use of k -sparse-eigencuts, and we have shown
498 a practical algorithm that can generate many cuts on the fly, using the connection of
499 the separation problem with SPCA. Additionally, we have seen in Section 4 that dense
500 cuts are usually quite strong, and therefore it is sensible to also consider them when
501 building a high-quality approximation. One needs to be careful though, since adding
502 too many dense cuts can impair the efficient solvability of the LPs, which was the
503 motivation of this work in the first place. In what follows, we show a computational
504 study which includes a way of balancing these forces.

505 **6. Computational experiments.** We present a computational study of the
 506 tradeoffs between strength and efficiency in eigenvector-based cuts that approximate
 507 the semidefinite constraint (1.2). To do this, we compare the performance of k -sparse-
 508 eigencuts and dense cuts in a pure cutting plane procedure.

509 **6.1. Implementation details.** All of our algorithms are implemented in C++.
 510 As in the experiments in Section 4, we use the Eigen library [32], version 3.3.7, for
 511 linear algebra computations and the LP solver is Gurobi 9.0 [33]. We use the C++
 512 Fusion API of MOSEK 9.2 [42] to solve SDP relaxations.

513 The k -sparse-eigencut oracle used in Algorithm 1 is the same described in Section
 514 5.2: we execute the Truncated Power Method (TPower) by Yuan and Zhang [67] after
 515 modifying the current iterate’s \tilde{M} as in Lemma 3.3.

516 All experiments were performed single-threaded in shared computing environ-
 517 ments. Most of the reported results used nodes from the Béluga cluster of Calcul
 518 Québec, in which each machine has two Intel Xeon Gold 6148 Skylake CPUs clocked
 519 at 2.4 GHz, and a variable amount of memory depending on the node. Some experi-
 520 ments with larger instances were run on a shared server with 512 GB of memory and
 521 thirty-two Intel Xeon Gold 6142 CPUs clocked at 2.6 GHz.

522 **6.1.1. Cutting plane algorithm.** We use a straightforward implementation of
 523 a cutting plane algorithm. We first solve (LP), the standard LP relaxation of (SDP)
 524 for bilinear terms, to obtain a solution (\hat{x}, \hat{X}) . If the corresponding \tilde{M} is not PSD,
 525 we use Algorithm 1 to find linear inequalities of the form (1.3) that are violated by
 526 \tilde{M} . We then resolve the LP updated with the new cuts, and we repeat the above
 527 process until we reach a terminating condition. We next specify which specific cuts
 528 are added in each iteration, how we obtain the LP optimal solution \tilde{M} , and what are
 529 the parameters and numerical tolerances used for our procedure.

530 *Cutting plane families.* We consider two cut families: k -sparse-eigencuts, which
 531 were extensively described throughout the paper, and dense cuts, which are the eigen-
 532 vectors corresponding to the negative eigenvalues of \tilde{M} . For dense cuts, we add one
 533 for every distinct negative eigenvalue.

534 The results of the enumeration experiments from Section 4 showed that even a
 535 few dense cuts are able to close a large amount of gap. On the other hand, sparse
 536 cuts are likely to yield faster LPs. Thus, with dense cuts, one may expect to have
 537 a fast improvement in the bound, followed by a tailing off as each iteration takes
 538 longer, as compared to sparse cuts for which more iterations might be possible, but
 539 each iteration may close less gap. We explore how these two phenomena interact with
 540 the following three algorithm types:

- 541 • **DENSE:** At each iteration, add all cuts obtained from the eigenvectors associated to
 542 negative eigenvalues of the matrix \tilde{M} .
- 543 • **SPARSE:** At each iteration, generate k -sparse-eigencuts through Algorithm 1; a more
 544 detailed description is given in Algorithm 2.
- 545 • **HYBRID:** This strategy begins with DENSE, switching to SPARSE when the LP solve
 546 time is at least the value of a parameter HYBRIDSWITCHINGTIME.

547 The motivation for HYBRID is to mix dense and sparse cuts, by focusing on rapid
 548 improvement in gap closed in the beginning through dense cuts, and then switching
 549 to sparse cuts to moderate the growth in the LP solution time.

550 *Solution to cut.* In an LP solver, one typically has three basic choices of opti-
 551 mization methods: *Simplex*, *Barrier*, and *Barrier with Crossover*. The first and third
 552 methods generate an extreme-point solution to the LP, while the second only guaran-
 553 tees a solution in the optimal face (which may or may not be a vertex). We conducted

Algorithm 2: SPARSE(LP, k): k -sparse-eigencuts

Input : Initial LP relaxation of (QCQP), e.g., (LP), and sparsity level $k \in \mathbb{N}$.
Parameters: NUMCUTSPERITER: maximum number of cuts per iteration;
TERMINATINGCONDITIONS: check when execution should terminate.
Output : A sequence of k -sparse-eigencuts $\{\tilde{w}_j\}_{j=0}^p$ such that $\|\tilde{w}_j\|_2 = 1, \|\tilde{w}_j\|_0 \leq k$
for all $j \in [p]$.

- 1 **Initialize:** $LP_1 \leftarrow LP, p \leftarrow 0, t \leftarrow 1$;
- 2 **while** TERMINATINGCONDITIONS *have not been met* **do**
- 3 Let \tilde{M} be an optimal solution to LP_t ;
- 4 Let $\{\tilde{w}_j^t\}_{j=0}^{p_t}$ be the output of SPARSEROUND(\tilde{M}, k), i.e., Algorithm 1, sorted in order
of decreasing violation with respect to \tilde{M} ;
- 5 $p'_t \leftarrow \min\{p_t, \text{NUMCUTSPERITER}\}$;
- 6 $LP_{t+1} \leftarrow LP_t$ with the addition of the first p'_t cuts of form (1.3) from $\{\tilde{w}_j^t\}_{j=0}^{p_t}$;
- 7 $\tilde{w}_{p+j} \leftarrow \tilde{w}_j^t$ for all $j \in [p'_t]$;
- 8 $p \leftarrow p + p'_t$;
- 9 $t \leftarrow t + 1$;
- 10 **end**
- 11 **return** $\{\tilde{w}_j\}_{j=0}^p$;

554 extensive preliminary experiments to determine which strategy to choose, and con-
555 cluded that *Barrier* always performs better for both sparse and dense cuts. It is not
556 surprising that the quality of the cuts is better in this case, as cutting off a point in
557 the interior of the optimal face can remove a larger portion of said face. However,
558 within a cutting plane algorithm, a sequence of LPs can often be solved substantially
559 faster with *Simplex*, using its warm start capabilities, than with *Barrier*. Yet in our
560 early experiments, the faster iterations afforded by *Simplex* led to significantly worse,
561 and ultimately slower, bound improvements than with *Barrier*.

562 *Cut management.* As discussed in Section 4, at each iteration of the cutting
563 plane algorithm, it is desirable to add a large number of k -sparse-eigencuts to ensure
564 sufficient progress. As accumulating an excessive number of inequalities can slow down
565 the LP solution time, especially when many of them become permanently inactive
566 after a few iterations, we implemented a simple cut management policy that is used
567 in all our experiments: a cut is removed from the pool if it is inactive for more than
568 one iteration.

6.1.2. Parameters.

569 *Sparsity.* We set a target sparsity level of $k = \lfloor 0.25(n + 1) \rfloor$. This implies that
570 all of the cuts that we add have fewer than 7% nonzero entries in the matrix space,
571 $M(x, X)$.

572 *Limits on cuts and supports.* We set NUMCUTSPERITER in Algorithm 2 to $5n$,
573 and MAXNUMSUPPORTS in Algorithm 1 to 100. We set HYBRIDSWITCHINGTIME,
574 the time limit for generating dense cuts in the HYBRID algorithm, to $\min\{10 \text{ seconds},$
575 $100 \text{ times the initial LP solve time}\}$.

576 *TPower initialization.* A significant detail in our implementation is the initial-
577 ization of TPower. The TPower method incorporates a truncation step (to achieve
578 sparsity) in the classical power method to compute the largest eigenvalue of a PSD
579 matrix. Since this algorithm is a heuristic for SPCA, the initialization plays a signifi-
580 cant role in the performance. We found that initializing TPower with the eigenvector
581 associated to the smallest eigenvalue generally leads to the best (and most consistent)
582 behavior.
583

584 *Tolerances.* We use a variety of tolerances throughout our procedure to ensure
 585 numerical stability. We say that the cut with respect to a vector v is violated by
 586 \tilde{M} if $v\tilde{M}v^\top < -\epsilon$, where we set $\epsilon = 10^{-7}$. An eigenvalue is considered negative
 587 if it is less than -10^{-6} . If the coefficient of a cut is less than 10^{-9} , we treat it
 588 as zero. To measure the progress of the algorithm, at each iteration, we compare
 589 the objective value, say z_{new} , to the previous iteration’s objective value, say z_{old} ; if
 590 $|z_{\text{new}} - z_{\text{old}}|/(|z_{\text{old}}| + \epsilon) < 10^{-5}$, we say that the objective has *stalled* for that iteration.
 591 For cut management, we track how often each cut is satisfied at equality; for this, we
 592 permit a tolerance of 10^{-3} .

593 *Terminating conditions.* We terminate the cutting plane algorithm (whether it
 594 is SPARSE, DENSE, or HYBRID) when one of the following conditions is met: (1) the
 595 time limit of 1 hour is reached, (2) no more cuts are added, (3) the objective did not
 596 sufficiently improve over the last 100 iterations.

597 **6.2. Instances.** We test three families of nonconvex quadratic problems, includ-
 598 ing several large instances with $n \geq 200$, i.e., with over 40,000 variables in the matrix
 599 space. In our representation of the problem, all variables (x, X) are present even if,
 600 for example, there is sparsity in (QCQP); solving the SDP using MOSEK (an interior-
 601 point optimizer) for these problems becomes extremely memory intensive, requiring
 602 35 GB or more for the larger instances, whereas the memory needed for the cutting
 603 plane algorithm remains relatively manageable. In addition to memory, the SDP also
 604 can take a significant amount of time; starting from $n = 200$, most SDP solution
 605 times are at least 1200 seconds, with some taking more than an hour. This suggests
 606 an advantage to the cutting plane approach in resource-constrained environments.

607 **BoxQP.** We consider the BoxQP family of QPs [15, 19, 59]. These are problems
 608 with nonconvex quadratic objective, and with box constraints $x \in [0, 1]^n$, that is,
 609 problems of the form

$$\begin{aligned} 610 \quad & \min_x x^\top Qx + c^\top x \\ 611 \quad & x \in [0, 1]^n. \end{aligned}$$

613 The experiments of Section 4 considered three of these instances. For BoxQP in-
 614 stances, it is well known that the SDP approach provides strong dual bounds to the
 615 nonconvex problem; we refer the reader to [19] for a semidefinite programming ap-
 616 proach to this class of instances and [21] for a completely positive approach to QPs.
 617 Recent papers [11, 64] have considered solving BoxQP with integer linear program-
 618 ming. We opted for not comparing with these approaches since, as we mentioned
 619 at the beginning, our goal is to mimic the effect of the semidefinite inequality (1.2).
 620 These alternative integer linear programming approaches cut through the SDP cone
 621 and derive valid inequalities for $X = xx^\top$ directly. This makes our approach rather
 622 complementary.

623 For this family, we consider all instances available at [https://github.com/sburer/](https://github.com/sburer/BoxQP_instances)
 624 [BoxQP_instances](https://github.com/sburer/BoxQP_instances), and we additionally generate our own larger instances, for $n = 200$
 625 and $n = 250$, using the same instance-generation code available in the link. In total,
 626 we conducted our experiments on 111 BoxQP instances, with $n \in [20, 250]$ and
 627 density $d \in [0.2, 1]$.⁴

628 **BIQ.** The BIQ instances are similar to the BoxQP instances: they only have a
 629 nonconvex quadratic objective function, although in this case the quadratic is homo-

⁴We define the *density* of a QCQP as the proportion of nonzero entries in $Q_0 = Q$.

630 geneous.⁵ The other difference is that the variables are restricted to be binary. In
 631 order to make these instances suitable for our setting, we reformulate each constraint
 632 $x_i \in \{0, 1\}$ as

$$633 \quad x_i(1 - x_i) = 0.$$

634 Therefore, the instances have the form

$$635 \quad \min_x x^\top Qx$$

$$636 \quad x_i(1 - x_i) = 0, \quad i \in [n]$$

$$637 \quad x \in [0, 1]^n,$$

639 which is a QCQP. The BIQ instances we consider are all the available instances on
 640 the Biq Mac library [63] with n up to 250. We do not consider the larger instances in
 641 the library due to memory limitations. We also exclude 14 instances for which none
 642 of the methods close any gap (e.g., if the initial solution is PSD).⁶ In addition, we
 643 remove one instance (bqp50-9) in which the objective value of the initial LP is within
 644 10^{-6} of the SDP optimal value. This leaves us with 135 instances from this family,
 645 with $n \in [20, 250]$ and $d \in [0.1, 1]$.

646 MAXCUT. We take the MAXCUT instances available in the Biq Mac library. It
 647 is well known that a MAXCUT instance over a graph $G = (V, E)$ with weights w_{ij} ,
 648 $\{i, j\} \in E$ can be formulated as

$$649 \quad \max_x \sum_{\{i,j\} \in E} w_{ij} \left(\frac{1}{2} - \frac{1}{2}x_i x_j \right)$$

$$650 \quad x_i^2 = 1, \quad i \in [n]$$

$$651 \quad x \in [-1, 1]^n,$$

653 which are also nonconvex QCQPs. Note that the bounds $x_i \in [-1, 1]$ are redundant, as
 654 the quadratic constraints imply $x_i \in \{-1, 1\}$. However, the initial relaxation benefits
 655 from having these bounds explicitly, which is why we left them in the formulation.

656 For this class of instances, it is known that the SDP relaxation provides a strong
 657 provable approximation of the problem [30]. We test all but 27 of the available
 658 instances in the Biq Mac library, pruning any with $n > 250$; this leaves us with 151
 659 instances, where $n \in [60, 225]$ and $d \in [0.1, 0.99]$.

660 **6.3. Computational results.** We use two comparison metrics. First, we com-
 661 pare the gap closed by all three methods, to assess relaxation *strength*. Second, we
 662 report the solution time of the last LP that is solved, to evaluate how *lightweight*
 663 the linear relaxations we are obtaining. Our most extensive analysis is for the BOXQP
 664 instances in Section 6.3.1; we curtail our discussion of the remaining families as the
 665 results lead to conceptually similar conclusions.

666 **6.3.1. BOXQP.** In Table 4, we summarize the performance of the SPARSE, DENSE,
 667 and HYBRID algorithms. We group instances by size, and the second column states
 668 the number of instances in each group. The next three columns give the percent gap
 669 closed, and the final three columns provide the time to solve the last LP.

670 The gap closed by the three methods is similar for smaller instances, and for
 671 $n \leq 125$ the relaxations obtained by all the methods are quite strong. The limit

⁵A *homogeneous* polynomial has the same degree for all nonzero terms.

⁶These are bqp50-1 through bqp50-8, gka1a, gka2a, gka3a, gka6c, gka7c, and gka8a.

TABLE 4

Results on 111 BoxQP instances for SPARSE, DENSE, and HYBRID. Results are averages over instances grouped by size, under a time limit of 1 hour.

Instance group	#	Gap closed (%)			Last LP time (s)		
		SPARSE	DENSE	HYBRID	SPARSE	DENSE	HYBRID
$n \in [20, 30]$	18	98.50	100.00	100.00	0.10	2.93	2.93
$n \in [40, 50]$	33	98.83	99.90	99.89	0.64	10.73	7.34
$n \in [60, 80]$	21	98.45	96.24	98.17	6.49	28.27	11.69
$n \in [90, 125]$	27	94.62	90.68	95.48	48.09	106.54	49.08
$n \in [200, 250]$	12	75.16	84.70	83.92	520.24	764.30	506.98

672 HYBRIDSWITCHINGTIME is never reached for the instances with $n < 40$, i.e., only
673 dense cuts are used, and DENSE and HYBRID have identical performance, while SPARSE
674 closes less gap, though with a corresponding final solving time for the LP of 0.1
675 seconds, compared to nearly 3 seconds for the other two methods.

676 For the instances with $n \in [40, 50]$, we see that HYBRID closes a little less of
677 the gap than DENSE, but with a 30% improvement in the LP solution time. Upon
678 a closer observation of this group, we observed that HYBRID encounters the limit
679 HYBRIDSWITCHINGTIME for 16 of the 33 instances, but it only adds sparse cuts for
680 4 of the instances. In the other 12 cases, HYBRID (via Algorithm 1) finds no sparse
681 cuts to add, terminating with fewer iterations, and a slightly lower gap closed, than
682 DENSE, but with a correspondingly lighter LP.

683 For the next two groups, $n \in [60, 80] \cup [90, 125]$, we see that sparse cuts, even in
684 isolation, outperform dense ones, whereas the gap closed by HYBRID is comparable to
685 that closed by SPARSE. As n grows, the importance of including dense cuts increases,
686 and eventually HYBRID dominates both SPARSE and DENSE. Curiously, this holds even
687 when accounting for the last LP time, showing that it is possible to have the best of
688 both worlds, i.e., both strength and speed, especially by combining dense and sparse
689 inequalities. This is more emphatic in the $n \in [90, 125]$ set: HYBRID closes nearly 5%
690 more of the gap than DENSE with an LP that solves over twice as quickly.

691 In the final and largest set of instances, in terms of gap closed, sparse cuts alone
692 now lag severely behind either of the procedures involving dense cuts, and the LP
693 time for HYBRID is the best of the three algorithms. We do observe a degradation in
694 the gap closed by HYBRID relative to DENSE, which we investigate further below.

695 In Figure 1, we focus on DENSE and HYBRID, and show a more detailed comparison
696 of the gap closed by each method. The dashed diagonal line indicates parity between
697 the two methods; instances above the line are ones in which HYBRID performs better,
698 and the opposite for instances below the line. Additionally, in this figure, we classify
699 instances according to their density.

700 The plot reinforces the message of Table 4, that a method combining sparse
701 and dense cuts has significantly stronger performance than one based on dense cuts
702 alone. We can further see that the improvement of HYBRID relative to DENSE gets
703 more pronounced on sparser instances, except for the ones in the lower left corner,
704 which we elaborate on next. There are seven instances in which at least one of the
705 two methods closes less than 82% of the gap: `spar125-025-1`, `spar200-025-*`, and
706 `spar250-025-*`. We focus on the three lying below the diagonal line, `spar250-025-1`
707 through `spar250-025-3`, for which HYBRID closes around 4% less gap than DENSE.




FIG. 1. Gap closed after 1 hour by HYBRID and DENSE algorithms on BOXQP instances, where the diagonal dashed line is a reference for equal performance. Most instances lie above this line, indicating that incorporating sparse cuts can close a substantially larger gap than a procedure with dense cuts alone. The size of the markers is proportional to n , and both methods tend to close less gap on larger instances within the time limit. Similarly, both methods perform worse on sparser instances, though HYBRID tends to have stronger performance relative to DENSE when an instance is sparser.

708 Unlike the situation for the rest of the family, HYBRID actually performs *fewer* iterations on these three instances. In these three cases, the reason is that an order of magnitude more sparse than dense cuts are produced, and ultimately the LP solution time for HYBRID tends to be slower than for DENSE, again in stark contrast to the rest of the family.

713 Our takeaway is that, although our heuristic parameter choices can be better tuned for larger BOXQP instances, overall HYBRID appears to be the better method on average, while SPARSE might be preferred for instances with $n \leq 80$, especially if LP solution time is weighed more heavily than gap closed, which can well be the case for a mixed-integer QCQP context.

718 We conclude this section with two plots detailing the progress of the three algorithms over the course of the hour. While this will be for only one instance, `spar125-025-1`, it exemplifies the relative behavior of the approaches that we observed during most of the experiments.

722 In Figure 2, the left panel shows the progress in gap closed throughout the hour, while the right panel shows the corresponding sequence of LP solution times. We see that DENSE suffers from tailing off, with a nearly flat curve after the first ten minutes of computation. Meanwhile, SPARSE eventually surpasses the gap closed of DENSE, though its rate of change is initially much slower. On the other hand, HYBRID, by design, captures the initial steep improvement achieved by dense cuts, after which it makes steady progress with the more efficient, but less aggressive, sparse cuts. Thus, if a shorter solution time is enforced, HYBRID would continue to dominate DENSE, whereas the same cannot be said for SPARSE compared to DENSE.

731 The right panel underscores the effect of the density of the cuts. The solution




FIG. 2. Progress of gap closed (left) and LP time (right) for BoxQP instance *spar125-025-1* for the three methods during the first hour.

732 times of DENSE are rapidly increasing, while for SPARSE, the LPs never take more than
 733 100 seconds to solve. HYBRID shows the expected middle-ground: the rapid initial gap
 734 closed comes at the expense of a nonnegligible initial increase of the LP solution times.
 735 However, switching to sparse cuts keeps the LP under control, while still being able
 736 to improve the gap closed. In fact, to test this further, we performed an extended
 737 test for this instance with a time limit of 1 day instead of 1 hour, at the end of which
 738 the LP solution times for SPARSE, HYBRID, and DENSE were 129.7, 127.1, and 375.8
 739 seconds, corresponding to a gap closed of 96.4, 96.3, and 88.1 percent, respectively.
 740 Hence, DENSE creates a sequence of increasingly heavy LPs with stagnating bound
 741 improvement, whereas SPARSE and HYBRID do not substantially slow down relative to
 742 the statistics after one hour and show sustained progress in the objective value.

743 *Comparison to Qualizza et al. [46].* We also evaluate HYBRID by implementing
 744 the best approach by Qualizza et al. [46], referred to as S2M, which was exten-
 745 sively tested on BOXQP instances. We modify the original S2M algorithm, in order
 746 to make a fair comparison with HYBRID, by enforcing the same rule based on
 747 HYBRIDSWITCHINGTIME used in HYBRID to determine when to start sparse cuts. In
 748 our experiments, HYBRID significantly outperforms S2M, closing substantially more
 749 gap within an hour, and HYBRID does increasingly better relative to S2M for larger
 750 instances. We refrain from adding detailed reports on this comparison since we imple-
 751 mented S2M ourselves; our tests with S2M qualitatively align with the results reported
 752 by Qualizza et al. [46], but subtle implementation details may differ.

753 **6.3.2. BIQ.** In Table 5, we summarize the performance of SPARSE, DENSE, and
 754 HYBRID for the 135 BIQ instances. The structure of the table is the same as Table 4.

755 We have similar conclusions as for the BOXQP instances. For the smallest set of
 756 BIQ instances, all three methods yield strong relaxations, with SPARSE closing about
 757 1% less gap than HYBRID, but having an 80% reduction in solving time, while HYBRID
 758 already comes with a 50% faster LP than DENSE, on average. In the next group,
 759 with $n = 100$, we see that HYBRID and SPARSE dominate DENSE, both closing around
 760 12.5% more of the gap on average while requiring only 34–38% of the time to solve
 761 the final LP compared to DENSE. This is a marked difference in gap closed, showing
 762 a substantially larger benefit to using sparse cuts in moderate-sized BIQ instances,
 763 compared to the more modest advantages we observed in the BOXQP family. Just
 764 as for the BOXQP instances, the relative performance, in terms of gap closed, of
 765 DENSE starts to improve again for larger n , while accordingly the quality of SPARSE

TABLE 5

Results on 135 BIQ instances for SPARSE, DENSE, and HYBRID. Results are averages over instances grouped by size, under a time limit of 1 hour.

Instance group	#	Gap closed (%)			Last LP time (s)		
		SPARSE	DENSE	HYBRID	SPARSE	DENSE	HYBRID
$n \in [20, 90]$	18	98.70	99.47	99.81	1.33	15.14	7.26
$n = 100$	31	94.92	82.53	95.00	31.33	91.35	34.82
$n \in [120, 150]$	41	90.18	89.35	92.61	125.87	262.56	132.43
$n \in [200, 250]$	45	54.72	65.72	64.06	479.61	830.75	519.96

TABLE 6

Results on 151 MAXCUT instances for SPARSE, DENSE, and HYBRID. Results are averages over instances grouped by size, under a time limit of 1 hour.

Instance group	#	Gap closed (%)			Last LP time (s)		
		SPARSE	DENSE	HYBRID	SPARSE	DENSE	HYBRID
$n = 60$	10	97.45	98.73	98.86	3.20	13.69	10.98
$n = 80$	30	93.61	93.43	96.65	18.59	47.48	24.07
$n = 100$	99	79.36	77.44	82.66	60.09	107.76	86.74
$n \in [150, 225]$	12	6.00	5.13	5.85	717.56	775.20	704.32

deteriorates, and HYBRID retains its status as a happy compromise of the two.

6.3.3. MAXCUT. Analogously to the other families, Table 6 summarizes our results for the 151 MAXCUT instances. The high-level trends remain the same as with the other two families, but there are a few notable differences for the largest group of instances. The gap closed is considerably less than for the other two families, with all methods closing only 5–6% of the gap. The reason is that the LP relaxation quickly becomes very heavy, and fewer than 15 iterations are able to be performed within the time limit.

In Figure 3, we take as a case study the MAXCUT instance `pm1s_100.1`, and we show the detailed evolution of all three methods over the course of the one hour time limit, with respect to gap closed (left panel) and LP solution time (right panel). For the time comparison in the right panel, the relative ordering of the algorithms is the same as for the BOXQP instance plotted in Figure 2, with SPARSE requiring the least amount of time throughout, followed by HYBRID, and then DENSE. As we did for `spar125-025-1`, here too we used an extended time limit of 1 day to observe the longer-run behavior: at the end of the day (not plotted), the LP solution was computed by for SPARSE, DENSE, and HYBRID in 58.5, 193.0, and 73.5 seconds, respectively.

However, the story for gap closed is different: while HYBRID continues to do well, picking up on the early momentum afforded by dense cuts and then continually increasing its relative advantage with respect to DENSE, the SPARSE algorithm seems to have much slower convergence than for the BOXQP setting. Indeed, after one day of computation time, SPARSE still trails DENSE in gap closed, with 92.6% of the gap closed by SPARSE, compared to 92.9% gap closed by DENSE. In the meantime, HYBRID improves from the one-hour mark gap closed of 89.6% to a final gap closed of 97.0%.




FIG. 3. Progress of gap closed (left) and LP time (right) for MAXCUT instance *pm1s_100.1* for the three methods over one hour.

791 **7. Concluding remarks.** The availability of strong, sparse outer approxima-
 792 tions of an SDP constraint is practically important, such as when solving QCQPs with
 793 LP-based spatial branch and bound. In this case, the need for such relaxations comes
 794 from the large number of LPs that need to be solved within that process and the
 795 target of quickly obtaining good bounds. A related, but distinct, motivation comes
 796 from convergence: without sparsity, a cutting plane approach may stall, either due to
 797 slower iterations or due to numerical issues.

798 This paper introduces a viable method to apply sparse eigenvector-based cutting
 799 planes to capture the strength of an SDP constraint in a linear programming context.
 800 While these families of cuts have been considered before, their empirical performance
 801 has been limited on one hand by the high density of eigenvector-based inequalities,
 802 causing linear programs that are too slow for practical purposes, and on the other
 803 hand by the relative weakness of sparsified versions of the dense cuts.

804 We first empirically justify, through our enumeration experiments, that sparse
 805 cuts tend to indeed be weak in isolation, while we observe that dense cuts are surpris-
 806 ingly strong when compared to all available sparse cuts. With this in mind, and via
 807 our novel connection between k -sparse-eigencut generation and the classical Sparse
 808 PCA problem, we develop HYBRID, an efficient separation routine for producing a mul-
 809 titude of strong and sparse eigenvector-based cuts. This method combines the initial
 810 strength of dense cuts and the steady progress that can be obtained when generating
 811 multiple sparse cuts using Algorithm 1. Our computational results indicate that we
 812 can successfully produce strong, lightweight linear relaxations; for QCQP instances
 813 with at most 100 variables and across three different benchmark families, within an
 814 hour of computation, we tend to close more than 90% of the gap between the initial
 815 LP relaxation and the SDP relaxation, representing a significant improvement with
 816 respect to either dense or sparse cuts alone. Further, our experiments indicate that
 817 a driving force for this relative improvement is that LPs with sparse cuts solve much
 818 faster: for example, in the last iteration, the LP with dense cuts can be two to three
 819 times slower, on average, than one with sparse cuts. Moreover, even with a much
 820 longer time limit, dense cuts often never attain the same gap closed as when sparse
 821 cuts are employed.

822 Future work involves further improving scalability. Right now, our handling of
 823 the variable matrix is straightforward, and all entries in X are created and stored.
 824 Avoiding this is a crucial step into scaling our approach into even larger instances and

825 to be able to properly embed it in spatial branch and bound.

826 **Acknowledgments.** This research was enabled in part by support provided
827 by Calcul Québec (calculquebec.ca) and Compute Canada (computecanda.ca). GM
828 would like to thank the Institute for Data Valorization (IVADO) for their support
829 through the Postdoctoral Fellow program, and to the Government of Chile for their
830 financial support through the FONDECYT grant number 11190515.

831 References.

- 832 [1] E. AMALDI, S. CONIGLIO, AND S. GUALANDI, *Coordinated cutting plane gen-*
833 *eration via multi-objective separation*, Math. Program., 143 (2014), pp. 87–110.
- 834 [2] K. M. ANSTREICHER, *Semidefinite programming versus the reformulation-*
835 *linearization technique for nonconvex quadratically constrained quadratic pro-*
836 *gramming*, J. Global. Optim., 43 (2009), pp. 471–484.
- 837 [3] M. ASTERIS, D. S. PAPAILIOPOULOS, AND G. N. KARYSTINOS, *Sparse principal*
838 *component of a rank-deficient matrix*, in 2011 IEEE International Symposium on
839 Information Theory Proceedings, IEEE, 2011, pp. 673–677.
- 840 [4] E. BALAS, *Intersection cuts—a new type of cutting planes for integer program-*
841 *ming*, Oper. Res., 19 (1971), pp. 19–39.
- 842 [5] R. BALTEAN-LUGOJAN, P. BONAMI, R. MISENER, AND A. TRAMONTANI, *Scoring*
843 *positive semidefinite cutting planes for quadratic optimization via trained*
844 *neural networks*. Working paper., 2019, [http://www.optimization-online.org/](http://www.optimization-online.org/DB_HTML/2018/11/6943.html)
845 [DB_HTML/2018/11/6943.html](http://www.optimization-online.org/DB_HTML/2018/11/6943.html).
- 846 [6] D. BIENSTOCK, C. CHEN, AND G. MUÑOZ, *Outer-product-free sets for polyno-*
847 *mial optimization and oracle-based cuts*, Math. Program., (2020), pp. 1–44.
- 848 [7] D. BIENSTOCK, C. CHEN, AND G. MUÑOZ, *Intersection cuts for polynomial*
849 *optimization*, in Integer Programming and Combinatorial Optimization, Springer
850 International Publishing, 2019, pp. 72–87.
- 851 [8] R. E. BIXBY, *Solving real-world linear programs: A decade and more of progress*,
852 Oper. Res., 50 (2002), pp. 3–15.
- 853 [9] G. BLEKHERMAN, S. S. DEY, M. MOLINARO, AND S. SUN, *Sparse PSD ap-*
854 *proximation of the PSD cone*, Mathematical Programming, (2020), pp. 1–24.
- 855 [10] G. BLEKHERMAN, S. S. DEY, K. SHU, AND S. SUN, *Hyperbolic relaxation of*
856 *\mathcal{K} -locally positive semidefinite matrices*, SIAM J. Optim., 32 (2022), pp. 470–
857 490, <https://doi.org/10.1137/20M1387407>.
- 858 [11] P. BONAMI, O. GÜNLÜK, AND J. LINDEROTH, *Globally solving nonconvex qua-*
859 *dratic programming problems with box constraints via integer programming meth-*
860 *ods*, Math. Program. Comput., 10 (2018), pp. 333–382.
- 861 [12] P. BONAMI, J. LINDEROTH, AND A. LODI, *Disjunctive cuts for mixed inte-*
862 *ger nonlinear programming problems*, Progress in Combinatorial Optimization,
863 (2011), pp. 521–541.
- 864 [13] P. BONAMI, A. LODI, J. SCHWEIGER, AND A. TRAMONTANI, *Solving quadratic*
865 *programming by cutting planes*, SIAM J. Optim., 29 (2019), pp. 1076–1105.
- 866 [14] E. BOROS, Y. CRAMA, AND P. L. HAMMER, *Chvátal cuts and odd cycle inequal-*
867 *ities in quadratic 0–1 optimization*, SIAM J. Discrete Math., 5 (1992), pp. 163–
868 177.
- 869 [15] S. BURER, *Optimizing a polyhedral-semidefinite relaxation of completely positive*
870 *programs*, Math. Program. Comput., 2 (2010), pp. 1–19.
- 871 [16] S. BURER, *A gentle, geometric introduction to copositive optimization*, Math.
872 Program., 151 (2015), pp. 89–116.

- 873 [17] S. BURER AND A. N. LETCHFORD, *On nonconvex quadratic programming with*
874 *box constraints*, SIAM J. Optim., 20 (2009), pp. 1073–1089.
- 875 [18] S. BURER AND R. D. MONTEIRO, *A nonlinear programming algorithm for solv-*
876 *ing semidefinite programs via low-rank factorization*, Math. Program., 95 (2003),
877 pp. 329–357.
- 878 [19] S. BURER AND D. VANDENBUSSCHE, *Globally solving box-constrained noncon-*
879 *convex quadratic programs with semidefinite-based finite branch-and-bound*, Comput.
880 Optim. Appl., 43 (2009), pp. 181–195.
- 881 [20] S. BURER AND Y. YE, *Exact semidefinite formulations for a class of (random*
882 *and non-random) nonconvex quadratic programs*, Math. Program., (2018), pp. 1–
883 17.
- 884 [21] J. CHEN AND S. BURER, *Globally solving nonconvex quadratic programming*
885 *problems via completely positive programming*, Math. Program. Comput., 4
886 (2012), pp. 33–52.
- 887 [22] A. CHMIELA, G. MUÑOZ, AND F. SERRANO, *On the implementation and*
888 *strengthening of intersection cuts for qcqps*, in Integer Programming and
889 Combinatorial Optimization, M. Singh and D. P. Williamson, eds., Cham,
890 2021, Springer International Publishing, pp. 134–147, [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-030-73879-2_10)
891 [978-3-030-73879-2_10](https://doi.org/10.1007/978-3-030-73879-2_10).
- 892 [23] S. S. DEY, B. KOCUK, AND A. SANTANA, *Convexifications of rank-one-based*
893 *substructures in QCQPs and applications to the pooling problem*, J. Global. Op-
894 tim., (2019), pp. 1–46.
- 895 [24] S. S. DEY, R. MAZUMDER, AND G. WANG, *A convex integer programming*
896 *approach for optimal sparse PCA*, arXiv preprint arXiv:1810.09062, (2018), [https:](https://arxiv.org/abs/1810.09062)
897 [//arxiv.org/abs/1810.09062](https://arxiv.org/abs/1810.09062). Working paper.
- 898 [25] S. S. DEY AND M. MOLINARO, *Theoretical challenges towards cutting-plane*
899 *selection*, Math. Program., 170 (2018), pp. 237–266.
- 900 [26] S. S. DEY, M. MOLINARO, AND Q. WANG, *Approximating polyhedra with sparse*
901 *inequalities*, Math. Program., 154 (2015), pp. 329–352.
- 902 [27] S. S. DEY, A. SANTANA, AND Y. WANG, *New SOCP relaxation and branching*
903 *rule for bipartite bilinear programs*, Optim. Eng., 20 (2019), pp. 307–336.
- 904 [28] M. FISCHETTI AND M. MONACI, *A branch-and-cut algorithm for mixed-integer*
905 *bilinear programming*, Eur. J. Oper. Res, (2019).
- 906 [29] M. FUKUDA, M. KOJIMA, K. MUROTA, AND K. NAKATA, *Exploiting sparsity*
907 *in semidefinite programming via matrix completion I: General framework*, SIAM
908 J. Optimiz., 11 (2001), pp. 647–674.
- 909 [30] M. X. GOEMANS AND D. P. WILLIAMSON, *Improved approximation algorithms*
910 *for maximum cut and satisfiability problems using semidefinite programming*, J.
911 ACM, 42 (1995), pp. 1115–1145.
- 912 [31] G. GRUBER, *On Semidefinite Programming and Applications in Combinatorial*
913 *Optimization*, PhD thesis, University of Klagenfurt, 2000.
- 914 [32] G. GUENNEBAUD, B. JACOB, ET AL., *Eigen v3*. <http://eigen.tuxfamily.org>,
915 2010.
- 916 [33] GUROBI OPTIMIZATION, LLC, *Gurobi optimizer reference manual*, 2020, [http:](http://www.gurobi.com)
917 [//www.gurobi.com](http://www.gurobi.com).
- 918 [34] R. HORST AND H. TUY, *Global Optimization*, Springer Berlin Heidelberg, 1996.
- 919 [35] M. JOURNÉE, Y. NESTEROV, P. RICHTÁRIK, AND R. SEPULCHRE, *Generalized*
920 *power method for sparse principal component analysis.*, J. Mach. Learn. Res., 11
921 (2010).
- 922 [36] J. B. LASSERRE, *Convergent SDP-relaxations in polynomial optimization with*

- 923 *sparsity*, SIAM J. Optimiz., 17 (2006), pp. 822–843.
- 924 [37] M. LAURENT, *Sum of squares, moment matrices and optimization over polyno-*
925 *mials*, in Emerging Applications of Algebraic Geometry, vol. 149 of IMA Vol.
926 Math. Appl., Springer, New York, 2009, pp. 157–270.
- 927 [38] J. LAVAEI AND S. H. LOW, *Zero duality gap in optimal power flow problem*,
928 IEEE T. Power Syst., 27 (2011), pp. 92–107.
- 929 [39] L. MACKEY, *Deflation methods for sparse PCA*, in Advances in Neural Informa-
930 tion Processing Systems, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou,
931 eds., vol. 21, Curran Associates, Inc., 2009, pp. 1017–1024.
- 932 [40] M. MAGDON-ISMAIL, *NP-hardness and inapproximability of sparse PCA*, Inform.
933 Process. Lett., 126 (2017), pp. 35–38.
- 934 [41] G. P. MCCORMICK, *Computability of global solutions to factorable nonconvex*
935 *programs: Part I – Convex underestimating problems*, Math. Program., 10 (1976),
936 pp. 147–175.
- 937 [42] MOSEK APS, *The MOSEK optimization toolbox for C++ manual. Version*
938 *9.2.21*, 2019, <https://docs.mosek.com/9.2/cxxfusion/index.html>.
- 939 [43] B. MÜLLER, F. SERRANO, AND A. M. GLEIXNER, *Using two-dimensional pro-*
940 *jections for stronger separation and propagation of bilinear terms*, SIAM J. Op-
941 tim., 30 (2020), pp. 1339–1365.
- 942 [44] G. MUÑOZ AND F. SERRANO, *Maximal quadratic-free sets*, in International
943 Conference on Integer Programming and Combinatorial Optimization, Springer,
944 2020, pp. 307–321.
- 945 [45] M. PADBERG, *The boolean quadric polytope: some characteristics, facets and*
946 *relatives*, Math. Program., 45 (1989), pp. 139–172.
- 947 [46] A. QUALIZZA, P. BELOTTI, AND F. MARGOT, *Linear programming relaxations*
948 *of quadratically constrained quadratic programs*, in Mixed Integer Nonlinear Pro-
949 gramming, vol. 154 of IMA Vol. Math. Appl., Springer, New York, 2012, pp. 407–
950 426.
- 951 [47] H. RAHMAN AND A. MAHAJAN, *Facets of a mixed-integer bilinear covering set*
952 *with bounds on variables*, J. Global. Optim., 74 (2019), pp. 417–442.
- 953 [48] M. V. RAMANA, *An algorithmic analysis of multiquadratic and semidefinite pro-*
954 *gramming problems*, PhD thesis, The Johns Hopkins University, 1994.
- 955 [49] J. K. REID, *A sparsity-exploiting variant of the Bartels-Golub decomposition for*
956 *linear programming bases*, Math. Program., 24 (1982), pp. 55–69.
- 957 [50] V. J. RODRIGUES DE SOUSA, M. F. ANJOS, AND S. LE DIGABEL, *Improving the*
958 *linear relaxation of maximum k-cut with semidefinite-based constraints*, EURO
959 Journal on Computational Optimization, 7 (2019), pp. 123–151.
- 960 [51] A. SANTANA AND S. S. DEY, *The convex hull of a quadratic constraint over a*
961 *polytope*, SIAM J. Optim., 30 (2020), pp. 2983–2997.
- 962 [52] A. SAXENA, P. BONAMI, AND J. LEE, *Convex relaxations of non-convex mixed*
963 *integer quadratically constrained programs: extended formulations*, Mathematical
964 programming, 124 (2010), pp. 383–411.
- 965 [53] H. D. SHERALI AND B. M. P. FRATICELLI, *Enhancing RLT relaxations via a*
966 *new class of semidefinite cuts*, J. Global. Optim., 22 (2002), pp. 233–261.
- 967 [54] N. Z. SHOR, *Quadratic optimization problems*, Sov. J. Comput. Syst. S.+ , 25
968 (1987), pp. 1–11.
- 969 [55] M. TAWARMALANI, J.-P. P. RICHARD, AND K. CHUNG, *Strong valid inequalities*
970 *for orthogonal disjunctions and bilinear covering sets*, Math. Program., 124
971 (2010), pp. 481–512.
- 972 [56] A. TILLMANN, *Computational aspects of compressed sensing*, PhD thesis, Tech-

- nischen Universität Darmstadt, 2014.
- [57] A. M. TILLMANN AND M. E. PFETSCH, *The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing*, IEEE Transactions on Information Theory, 60 (2013), pp. 1248–1259.
- [58] H. TUY, *Concave programming with linear constraints*, in Doklady Akademii Nauk, vol. 159, Russian Academy of Sciences, 1964, pp. 32–35.
- [59] D. VANDENBUSSCHE AND G. NEMHAUSER, *A branch-and-cut algorithm for non-convex quadratic programs with box constraints*, Math. Program., 102 (2005), pp. 559–575.
- [60] M. J. WAINWRIGHT AND M. I. JORDAN, *Treewidth-based conditions for exactness of the Sherali-Adams and Lasserre relaxations*, Tech. Report 671, University of California, September 2004.
- [61] M. WALTER, *Sparsity of lift-and-project cutting planes*, in Operations Research Proceedings 2012, Springer, 2014, pp. 9–14.
- [62] A. L. WANG AND F. KILINÇ-KARZAN, *On the tightness of SDP relaxations of QCQPs*, Math. Program., (2021).
- [63] A. WIEGELE, *Biq Mac Library*, 2007, <http://biqmac.uni-klu.ac.at/biqmaclib.html>. Accessed August 14, 2020.
- [64] W. XIA, J. C. VERA, AND L. F. ZULUAGA, *Globally solving nonconvex quadratic programs via linear integer programming techniques*, INFORMS J. Comput., (2019).
- [65] Y. YAJIMA AND T. FUJIE, *A polyhedral approach for nonconvex quadratic programming problems with box constraints*, J. Global. Optim., 13 (1998), pp. 151–170.
- [66] Y. YE, *Approximating quadratic programming with bound constraints*, Math. Program., 84 (1997), pp. 219–226.
- [67] X.-T. YUAN AND T. ZHANG, *Truncated power method for sparse eigenvalue problems*, J. Mach. Learn. Res., 14 (2013), pp. 899–925.
- [68] A. YURTSEVER, J. A. TROPP, O. FERCOQ, M. UDELL, AND V. CEVHER, *Scalable semidefinite programming*, SIAM Journal on Mathematics of Data Science, 3 (2021), pp. 171–200.